OPTICAL NETWORK TESTBEDS WORKSHOP 2 (ONT2): Full Report

Table of Contents

Executive Summary

1. Introduction

A new generation of optical networking services and technologies is rapidly changing the world of communications. National and international networks are implementing optical services to supplement traditional packet-routed services. On September 12-14, 2005, the Optical Network Testbeds Workshop 2 (ONT2), an invitation-only forum hosted by the NASA Research and Engineering Network (NREN) and co-sponsored by the Department of Energy (DOE), was held at NASA Ames Research Center in Mountain View, California. The aim of ONT2 was to help the Federal Large Scale Networking Coordination Group (LSN) and its Joint Engineering Team (JET) coordinate testbed and network roadmaps describing agency and partner organization views and activities for moving toward next-generation communication services based on leading-edge optical networks in the 3-5 year timeframe.

ONT2 was conceived and organized as a sequel to the first Optical Network Testbeds Workshop (ONT1, August 2004, www.nren.nasa.gov/workshop7). ONT1 resulted in a series of recommendations to LSN. ONT2 was designed to move beyond recommendations to agree on a series of "actionable objectives" that would proactively help federal and partner optical network testbeds and advanced research and education (R&E) networks to begin incorporating technologies and services representing the next generation of advanced optical networks in the next 1-3 years.

Participants in ONT2 included representatives from:

- Innovative prototype networks (Panel A)
- Basic optical network research testbeds (Panel B)
- Production R&D networks, including "JETnets," selected regional optical networks (RONs), and international R&D networks (Panels C and D)
- Network researchers (Panel E)
- Commercial network technology and service providers (Panel F)

- Senior engineering and R&D managers from LSN agencies and partner organizations

The overall goal of ONT2 was to identify and coordinate short- and medium-term activities and milestones for researching, developing, identifying, evaluating, and implementing the services, technologies, and interoperability mechanisms required. The roadmaps were formulated and presented not so much to reconcile the roadmaps with each other, but rather to provide a means to compare the major ongoing and planned optical networking activities in the R&E community, organized by categories of activities and communities of interest.

In addition, a 5-15 year network research perspective was provided by Panel E, which presented a report on two recent National Science Foundation workshops that examined long-term research goals and directions for optical internetworking. The report, "Mapping a Future for Optical Networking and Communications," by Dan Blumenthal (panel presenter), John Bowers, and Craig Partridge, is available at http://www.ocpn.ece.ucsb.edu/workshop/nsf-opt-rpt-final_R4.pdf

Finally, industry perspectives on forthcoming optical networking technologies and services were presented in Panel F by representatives of companies from the optical technologies and network services industries.

The presentations and roadmaps given at the workshop are available on the ONT2 website (www.nren.nasa.gov/workshop8).

2. Workshop Results and Next Steps

ONT-2 developed a series of specific "actionable objectives" or frameworks for community action as ongoing mechanisms to ensure measurable progress toward the defined goals, and to ensure that barriers to this progress are addressed. Key examples of these actionable objectives are given in the next section of this executive summary.

Although developing these objectives or frameworks was a goal of ONT2, the workshop was not intended to fully develop the required processes needed to achieve these objectives, but rather to point to "completion activities" that the community needs to undertake subsequent to the workshop. The workshop sponsors and organizers recognized that there are several existing related processes and forums, such as standards bodies, the Global Grid Forum, and especially the Global Lambda Integrated Facility (GLIF) consortium, which are already active in this area, but need further coordinated participation and support from the LSN and partner R&E networking communities.

The ONT2 workshop report consists of several parts including this Executive Summary and summary reports from the panels and breakout sessions. These documents have been approved by LSN and are available on the ONT2 website: www.nren.nasa.gov/workshop8.

The global next step is to promote international coordination in this area. To this end, plans are now well underway by the Japanese NICT Institute (National Institute of Information and Communications Technology), in cooperation with related US organizations coordinated by LSN, to organize a subsequent ONT3 invitation-only workshop that will focus on international optical network testbeds coordination, to be held in Tokyo in September 2006, in conjunction with the GLIF annual meeting.

3. Key "Actionable Objectives" for the JETnet Community

Actionable objectives agreed upon at ONT2 include the following, which were selected as key examples from the individual ONT2 Panel A, Panel C/D, and Service Definitions Summary Reports. Please read those individual summary reports in their entirety for the complete list of actionable objectives agreed upon at ONT2, along with a much deeper understanding of the themes and issues involved.

3.1 Working Toward Multi-domain Interoperability

It's too early to work top-down—making something work across two or three domains is enough of a challenge—first at the control plane layer, then at higher layers. (Data Plane is largely a solved problem, being worked within GLIF for global interoperability.)

Multi-domain Control Plane is a multiyear problem, and will be solved in an evolutionary way by bilateral inter-testbed and inter-regional efforts, but all testbeds should be talking with each other *globally*, through the GLIF and GGF forums.

Senior network engineers should **work at iGRID and within GLIF** to agree on common sets of **functionalities for Open Exchange Points.**

JETnets are invited to participate with the **ESnet OSCARS project and the Internet2 BRUW project**, in coordination with GEANT2 and CANARIE, on **working together and within GLIF in the intra- and inter-domain dynamic provisioning and routing area**.

3. 2 Promoting Distributed Applications and Use of Layer 1/Layer 2 Services

The following are selected examples of actionable objectives in this area that were agreed in the Panel A breakout sessions:

Testbeds should **make light paths available to selected applications**, and work with the applications on how best to begin using lambdas.

GIG-EF and DRAGON are already working together on Haystack, and will share their approaches and experiences with the community.

HOPI and USN agreed to get together in the next few months **to identify an application on which they can collaborate.**

GIG-EF and USN are both interested in distributed file systems like LUSTRE, and will look into **collaborating on support for distributed applications over LUSTRE.**

3.3 Agreeing on Common Service Definitions

Testbeds like **DRAGON should identify key services** that they need to offer their users, **then develop engineering specifications** (including cost/price as appropriate) followed by encodings in appropriate formal languages, e.g., BNF or XML.

The community should get started right away to **define and adopt an initial set of basic light path service definitions** by working within the GLIF framework.

The community should have a number of **Open Exchange Points up and running with services based on agreed common service definitions by September 2006**, and review progress in this area at ONT3.

4. Key "Next Steps" for the Research Testbed Community

The next steps agreed upon at ONT2 by the research testbed community are the following, taken directly from the Panel B report. Please read that report in its entirety to get a good understanding of the discussion of research testbed themes and issues that led to agreement on the following "next steps" for the community.

4. 1 Research Testbed Areas of Common Interest

The optical networking research testbed community should **identify opportunities for collaboration with the GLIF community in the area of service definitions**.

In the area of next-generation management and control planes, the potential for innovative cooperative research projects among international testbeds, working with the GLIF community, represents an important opportunity.

4.2 Need for Large-Scale Multi-Year Research Testbed

The research community should move toward **support for a large-scale optical research testbed that would involve research participants from other (related) investigative knowledge domains**.

The research community should undertake an effort to define the **appropriate funding models that will support key research objectives over a multi-year timeframe**.

5. Vision Themes and Enablers

Throughout the workshop, views of an "overall vision" of future networking were articulated and discussed during both ONT2 panel discussions and breakout sessions. Perhaps the community needs to agree on such an overall vision, or more likely, the community needs only to move together towards the realization of whatever the future may hold in this area. Several "vision" themes were heard over and over again in several panels—these vision themes and related enablers and activities are summarized in this section of the executive summary. They are presented for consideration by the community, and provide further context for the workshop report.

5.1 Entirely New Kind of Network for Real Applications

The advent of Layer 1/Layer 2 networking as a viable technology for solving performance problems (of several kinds) in the wide area is going to create an entirely new kind of network over the next few years. It is important that the LSN and partner communities—and specifically the JET—get out in front to help coordinate these research prototyping testbed networks as they start moving toward solving the performance problems of real application scientists doing real agency work.

5.2 The "Vision"

In its simplest and most compelling "overall vision," Layer 1/Layer 2 networking adds an "underlay" to existing internetworking, compared with multiprotocol label switching (MPLS) which adds an "overlay." Taken together, these new, layered services will create a rich, new network architecture. Direct light paths will allow incoming high-performance connections to bypass Internet routers and firewalls (a.k.a., "bottlenecks") and connect directly to campus high-performance data sources, supercomputer servers, and visualization users. New protocols will provide for higher speeds over longer distances, over either a single light path or an Optical Private Network (OPN). Such networks and protocols will be user-controlled to an unprecedented degree, enabling applications to call upon and tie together high-performance services deep within the network.

5.3 Immediate Need to Buy Dark Fiber

The R&D needed to pursue this vision in the next few years will require major access to dark fiber, whether agency-owned or shared with other agency, national, or regional networks. The window of opportunity is closing rapidly for being able to buy sufficient dark fiber to support the needed long-term R&D at today's reasonable prices. In the U.S., the National LambdaRail (NLR) and the Regional Optical Networks (RONs) have gone a long way toward meeting this need, however, it was noted that agencies and partner organizations should get on board with these opportunities immediately, as again, the window is closing rapidly.

5.4 Changes Needed at Campus Level

For both federal laboratories and university campuses, network managers will have to shift their thinking, becoming network *infrastructure* managers instead. In the optical network generation, there will be several networks coming into the campus, some of which will not be managed at the campus level. Every institution will need to build a fiber infrastructure to its RON to enable these multiple networks to enter and leave the campus.

5.5 GLIF

Globally, a very important coordination activity in this area is being provided by (GLIF) (www.glif.is). GLIF is an international virtual organization that promotes the paradigm of lambda networking. GLIF provides lambdas internationally as an integrated facility to support data-intensive scientific research, and supports middleware development for lambda networking. GLIF brings together senior networking engineers to develop an international infrastructure by identifying equipment, connection requirements, and needed engineering functions and services. The GLIF participants are national research and education networks (NRENs), consortia, and institutions working with lambdas. Administrative support is provided by the Trans-European Research and Education Networking Association (TERENA) (http://www.terena.nl/) with financial support from sponsoring organizations.

5.6 GENI

On the research side, a very important long-term program aimed at developing and realizing the “vision” is the NSF Global Environment for Networking Investigations (GENI) (www.nsf.gov/cise/geni). The aim of GENI is to explore new networking capabilities that will advance science and stimulate innovation and economic growth. The GENI initiative includes a research program and a global experimental facility designed to explore new network architectures at scale. This research initiative will explore creating new core functionality, developing enhanced capabilities, deploying and validating new architectures, building higher-level service abstractions, building new services and applications, and developing new network architecture designs. The GENI facility will extend beyond existing testbeds to create an experimental infrastructure capable of supporting the ambitious research goals of the GENI initiative. NSF envisions GENI as contributing to, and promoting a broad community effort that engages other agencies, countries, and corporate entities.

Summary Report From Panel A

("Breadth First" Testbeds and Related Breakout Group Discussions)*

Panelists:

UltraScience Net (USN), Bill Wing (Panel A Co-chair)
HOPI, Rick Summerhill (Panel A Co-chair)
National LambdaRail (NLR), John Jamison
Global Information Grid Evaluation Facility (GIG-EF), Hank Dardy
TransLight/GLIF, Kees Neggers
CA*net4, Bill St. Arnaud

1. Overview

The ONT2 Panel A breakout sessions began by listening to each of the Panel A testbeds express their primary themes as described in their Panel A presentations and roadmaps, then combined these themes into three overall themes with associated issues:

Theme 1: Interoperability

Theme 2: What Do Users Need? What Do We Provide?

Theme 3: Open Exchange Points

The breakout discussion on each of these themes is summarized in the following sections, together with the "Actionable Objectives" that were agreed upon for Panel A participating organizations to pursue following the workshop. Details of the development of these three themes are contained in the appendix to this Panel A summary report.

2. Theme 1: "Interoperability"

2.1 Summary of Panel A Testbeds Discussion on Theme 1: "Interoperability"

2. 1.1 HOPI

OC192 to Europe, so have connections to London, Amsterdam.
Status: Moving data; Control plane is in place but not yet opened up to the outside world.
Initial approach: Provide support for the UltraLight project (physics applications), need interoperability with DRAGON, CHEETAH, regional pieces.

2.1.2 UltraScience Net

Interoperability with CHEETAH is mandated. Overall approach in three parts:

- Provide control plane user interface.
- Interoperate with CHEETAH.
- Support CHEETAH interoperability with DRAGON and HOPI.

* This panel pertains to optical network testbeds whose focus is on engineering evaluation of best technologies and services to establish and serve a large community.

2.1.3 GIG-EF

Working at the research level with Jerry Sobieski et al. on MAX, ATDnet (riding same fibers), DRAGON, BOSSnet to Haystack. Overall approach in two parts:

- Can now switch wavelengths between these networks (using Movaz gear) and to NRL (physical layer data plane and control plane).
- Next, get Session Initiation Protocol working between networks, working with Terry Gibbons, Tom Lehman.

2.1.4 SURFnet

Working primarily through the Global Lambda Integrated Facility (GLIF). Overall approach in three parts:

- Interoperability is working at Data plane Layer 1 through Netherlands (NetherLight), Chicago (StarLight), Seattle (Pacific Wave), New York (Manhattan Landing). GÉANT2 is also fully compatible at Layer 1. So Layer 1 is a given, solved by working with GLIF.
- Next need is to get agreement on global approaches to control plane interoperability (within each cloud, just use own vendor's approach, so the problem is between clouds). Too early to decide on global standards, but can begin working on issues—main thing is to make sure there is a **global approach**, not a number of national or regional approaches. GLIF and the Global Grid Forum (GGF) are the appropriate forums for discussions. Gigi Karmous-Edwards is working within GLIF on control plane issues.
- Higher layers: Use Ethernet framing, but otherwise no defined international connections at Layer 2; just bring your lambdas, e.g., to MAN LAN Ethernet switch, set up bilateral arrangements. For higher layers, could use SIP, e.g., to request Layer 2 and 3 services for video applications.

2.2 Actionable Objectives for "Interoperability"

Data Plane is getting mature. Work with GLIF for global interoperability.

Multi-domain Control Plane is a multiyear problem—will be solved in an evolutionary way by bilateral inter-testbed and inter-regional efforts, but all testbeds should be talking with each other **globally**, through the GLIF and GGF forums.[**]

Service Plane. Follow the discussions and approaches on Common Services.

Applications. Testbeds should **make light paths available to selected applications**, and work with the applications on how best to begin using lambdas.

3. Theme 2: "What Do Users Need? What Do We Provide?"

3.1 Summary of Panel A Testbeds Discussion on Theme 2: "What Do Users Need? What Do We Provide?"

3.1.1. HOPI

Testbed Support Center basically has three pieces:

- The engineering piece, provided by the Indiana University Network Operations Center, now working.
- The HOPI control plane, now in beta.
- An applications focus on both: (a) what is needed to support existing applications, and (b) what new thinking is needed to develop new applications over different networking paradigms? North Carolina is building sets of equivalence classes of networking services for support of new applications. Aiming to have all three pieces working together and being used by new applications by September 2006.

3.1.2 GIG-EF

Strategy for encouraging and supporting new applications has basically three parts:

- Building applications that can decide whether to go standard Grid versus to InfiniBand, e.g., visualization.
- Moving InfiniBand outside the computer room to the WAN.
- Providing distributed file systems on the desktop.

[**] Following the workshop, one LSN member provided the following important "watch-outs" concerning this item, which should be taken into account by the community going forward: As the community migrates toward global control plane, there must be limits on span of control. Designers, implementers, and operators need to consider how the system would function on those days when things were going badly and cascading faults were propagating through the network. "Diversity"—of many types and at many points—would be a crucial tool for managing the propagation of faults and the potentially crippling effects of network failures. During the Panel A breakout discussion, it was noted that having potentially different control planes for different domains would be a possibility, as long as they would interoperate with each other. Eventually, standards could evolve toward a single control plane, much like BGP provides today in the IP routing world, but that would take time, of course.

3.1.3 UltraScience Net (USN)

Four key applications:

- Remote control of high-speed optical confocal microscope, which generates high-resolution imagery with requirements for low latency and low jitter. Application streams image data for analysis, which is rendered into video products, then streamed to users.
- Lambda Station: High Energy Physics application developed as part of DOE Large Hadron Collider effort—uses intelligent data steering to reserve USN bandwidth, and then take optimal advantage of it.
- Dedicated data and control-channel-based visualization application being developed for the Terascale Supernova Initiative. Application takes advantage of USN's ability to create parallel, independent dedicated channels at specified bandwidths between the same two endpoints.
- LUSTRE: geographically distributed file system for supercomputers (Cray X2 and XTE). LUSTRE is an open source file system being developed at several centers, designed to operate as a distributed system from its inception. In this context, USN is serving as a testbed for extending it to the wide area, taking advantage of the fact that at Layer-2, USN can operate without router-generated session delays.

Many similarities to the requirements and strategies as GIG-EF, should collaborate.

3. 2 General Discussion

All testbeds are interested in open source middleware for support of new applications, not proprietary middleware.

Applications hardware is getting cheaper and cheaper, e.g.:

- Tile displays now available for $20K can consume five Gigabit/s (Gbps), need a dedicated lambda to run transparently.
- Single low-cost end devices can now generate or consume up to eight Gbps.
- High-end servers available from Dell and IBM for $10K can burst at 30 Gbps.

What efforts are going on to develop new APIs? This is where Service Oriented Architectures (SOAs) fit in, providing clean APIs at key levels to the community. We need to do this together. Jerry Sobieski wants to develop support for "application specific topology," where an application request can quickly set up a target topology at multiple layers. Working on this together with partners, discussing and exchanging notes, comparing efforts, would eventually lead to the multilayer, multi-domain interoperability for this type of application support.

All agencies have their own facilities and users plus university researchers. So interoperability of applications support for the USN, HOPI, and GIG-EF communities makes sense. "We need to get the collaboration within the Federal government back on track." DARPA and NSF used to force some of this by having joint PI meetings, but this is not happening today. USN is funded by both DOE and NSF, the researchers are at universities. GIG-EF and DRAGON share some of the same channels to universities. NASA also needs to get more involved with this community—excellent relationships with DOD networks years ago have disintegrated in recent years.

3.3 Actionable Objectives for "What Do Users Need? What Do We Provide?"

- **It's too early to work top down**. Making something work across two or three domains is enough of a challenge, first at the control plane layer, then at higher layers.
- **Use opportunities like iGrid and SC05 events** to move forward incrementally, but trying to work together rather than in islands. Discuss, share, and adapt to the better approaches and practices.
- **GIG-EF and DRAGON are already working together on Haystack**, and will share their approaches and experiences with the community.
- **HOPI and USN agreed** to get together in the next few months **to identify an application on which they can collaborate.**
- **GIG-EF and USN** are both interested in distributed file systems like LUSTRE, and will look into **collaborating on support for distributed applications over LUSTRE.**

4. Theme 4: "Open Exchange Points"

4.1 Summary of Panel A Testbeds Discussion on Theme 3: "Open Exchange Points"

"Open Exchange Points" (OEPs) are already happening at Layer 1 ("Bring me your lambdas"):

- Manhattan Landing (MAN LAN) in New York
- NetherLight in Amsterdam
- Pacific Wave in Seattle
- StarLight in Chicago.
- MAX has the NGIX-East in the Washington DC area, which is switching inter-testbed traffic at Layer 1 and soon may be advertised as a Layer 1 OEP.
- SOX in Atlanta and AMPATH in Miami are also currently working on this.

With the right framing (Ethernet), such OEPs automatically support bilateral agreements at Layer 2. OEPs typically provide IP/BGP support for Layer 3 as well.

There is interest in the U.S. in developing distributed OEPs (e.g., Pacific Wave has Ethernet/HDXCs in Los Angeles and Seattle interconnected with a dedicated E-wave(s)).

4.2 Objectives for "Open Exchange Points"

- Senior network engineers should **work at iGRID and within GLIF** to agree on **common sets of functionalities for OEPs**.
- Sponsoring organizations should **"Think globally, act locally,"** to provide OEPs for the community at key points of global interconnection.
- **Issues to work on:** Middleware, AAA to decide peering, call blocking.

APPENDIX: Panel A Primary Themes from USN, HOPI, GIG-EF, SURFnet, CA*net4

CA*net 4 Primary Themes
- How do we provide users more control to manage, deploy, and use our networks?

HOPI Primary Themes
- Operations that serve all layers of the protocol stack
- Providing support for new applications
- Interoperability in this new paradigm

DOE USN Primary Themes
- Triangle of themes all interrelated: Interoperability, Services, Security

SURFnet Primary Themes
- Access to dark fiber, backbone, and customer
- Inter-domain issues
- Protocols
- Support
- End-to-end across globe on different nets
- Exchange points

GIG-EF Primary Themes
- Native IPv6 everywhere
- Build out control plane
- SIP
- Web services
- End-to-end QoS/QoP for high-end streams

Panel A Combined Themes–Individual Testbed Themes Combined Into Three Main Themes With Associated Issues

Interoperability

- Optical Transport
- Data Plane
- Control Plane
- Management Plane
- Routing Plane
- Service Plane
- How to provide fault isolation

What Do Users Need? What Do We Provide?

- Services?
- SOA
- What tools can be provided so users can manage/provision their own networks?
- Stimulate new applications on these networks (middleware?)

Open Exchange Points

- Need to be well-defined at lower layers
- Need framework/standards
- Dynamic provisioning
- Identify and mitigate risks
- Tools for provisioning

Summary Report From ONT2 Panel B

("Depth First" Testbeds and Related Breakout Groups)*

Panelists:

NSF Guru Parulkar (Panel B Co-chair)
OptIPuter, Tom DeFanti (Panel B Co-chair)
DRAGON, Tom Lehman, Jerry Sobieski
CHEETAH, Ibrahim Habib
OMNInet, Joe Mambretti
Japan Gigabit Network 2 (JGN2), Kunihiro Kato
VIOLA, Peter Kaufmann

1. Overview

Panel B was organized to present recent initiatives related to optical network testbeds focused on basic research addressing specific challenges "in-depth." Presenters provided an overview of the objectives and results of these research initiatives. Through five-year roadmaps, they identified long- and short-term goals, as well as progress toward developing next generation optical services, technologies, and interoperability mechanisms. To place these five-year roadmaps within a larger context, the Panel was organized to be complementary to the subsequent Panel E, which presented a 5-15 year perspective as reflected in the results of two recent National Science Foundation workshops. In a breakout session following the presentations and Q&A, panel members formulated the following key issues and themes.

2. Key Issues and Themes

Panel members stressed the importance of having common optical network services definitions across multiple optical network domains. It should be possible to interoperate and concatenate these services, which, in part means that they should have repeatable, known, consistent characteristics. This is a requirement for inter-network operability. It was noted that other organizations were addressing these issues, and that the work of those organizations, especially the Global Lambda Integrated Facility (GLIF), should be leveraged.

Another major theme was the recognition that the research community requires an experimental optical network, focused on long-term versus short-term goals that scales both at the infrastructure and at the community level. Such a testbed would be defined by multiple characteristics. For example, it would be oriented to solving challenging problems based on difficult-to-develop, but promising technologies with the potential for

* This panel pertains to optical network testbeds with a focus on research exploration of particular technologies.

being highly disruptive. The testbed would investigate many advanced optical devices, subsystems, systems, and methods currently being developed in labs.

The testbed would also involve larger communities of interest such as leading-edge science communities that traditionally drive leading-edge technology. This activity should have a "teambuilding" process that scales to multiple research communities. To advance this goal, partnerships should be formulated with advanced technology communities, such as those developing the advanced technology components of the TeraGrid and datagrids. This testbed would involve Federal agencies, research centers, and international testbeds. In part, this goal requires developing techniques for bringing these new technologies, such as light path services, into individual labs. Current campus networks comprise a barrier to advanced optical services.

Another key theme was the recognition of the importance of emerging new architectures for distributed management and control planes. Here also, it was recommended that the activities of the GLIF could be examined as a potential vehicle to leverage such activities. Basic research issues include: developing common definitions; inter-domain and application policy systems and instantiation methods at multiple levels (e.g., network, operations, business); interoperability methods; and measurement (e.g., security, data integrity, performance, validation).

To accomplish these network research goals, persistent funding over time must be provided both for the basic research and for regularly refreshing the optical testbed technology. Also, the Panel noted that research funding should be established specifically for interdisciplinary projects. In part, this approach requires changing existing policy at some agencies, which are oriented toward support for existing disciplinary areas. In addition, it was noted that the current funding for optical testbed research does not include support for innovative research projects among testbeds. Consequently, such projects are limited and ad hoc.

3. Next Steps

- The optical networking testbed research community can benefit in the area of services definition by leveraging existing efforts, such as those undertaken by the GLIF community. Opportunity for collaborative efforts with that community should be identified.
- A planning effort should be undertaken, by optical networking researchers to define the objectives, processes, operations, and support models for a large-scale optical research testbed that would involve research participants from other (related) investigative knowledge domains.
- The optical networking testbed research community should also explore the potential for leveraging existing efforts related to the design and development of next generation management and control planes, especially among the GLIF community. The potential for innovative, cooperative research projects among international testbeds represents an important opportunity.
- The research community should undertake an effort to define the appropriate funding models that will support their key objectives over a multi-year timeframe.

Summary Report From ONT2 Panels C and D

(Panel C = "Already Married" JETnets[*] and Panel D = "Still Looking" JETnets[**])

Panelists:

Panel C:
Abilene (with HOPI), Steve Cotter (Panel C Co-chair)
DREN (with GIG-EF), Phil Dykstra (Panel C Co-chair)
ESnet (with UltraScience Net), Chin Guok
GEANT, Maartin Büchli
SURFnet, Kees Neggers

Panel D:
NASA, Kevin Jones (Panel D Co-chair)
NCREN and RONs, Mark Johnson (Panel D Co-chair)
USGS and EROS Data Center, Mike Rechtenbaugh
NOAA, Alex Hsia
NIH/NLM, Jules Aronson
GLORIAD, Greg Cole (via video link)

1. Overview

Panels C and D (consisting primarily of the JETnets) held breakout sessions to identify key themes and issues associated with the development and implementation of optical networks. In addition to the issue of Common Service Definitions, which was shared with the other panels, Panel C and D breakout session participants identified four themes of interest:

Service Provisioning: The process of securely allocating or sharing network resources—this includes instantiating network services within or across autonomous domains or vendors.

Operations and Management: The "care and feeding" of users and the network while maximizing service offerings.

Cost Models: Necessary to ensure the network's financial sustainability.

"MorphNet" Architecture: "Provide support for multiple, concurrent, multi-layer views of the network for the applications and the network researchers" (see MorphNet URL in Section 6).

---

[*] This panel pertains to networks primarily funded by organizations with own in-house optical network testbeds.
[**] This panel pertains to networks, which primarily want to position themselves downstream from non-own optical network testbed technologies, systems, and service developments.

2. Actionable Objectives for JETnets

During the breakout session discussions, one specific actionable objective was identified in the area of Service Provisioning: JETnets are invited to participate with **the ESnet OSCARS project and the Internet2 BRUW project**, in coordination with GEANT2 and CANARIE, to **work together in the intra- and inter-domain dynamic provisioning and routing areas**.

In addition to this specific actionable objective, several synergies and potential collaborations were identified in all four areas of interest. These items are woven into the following. To foster optical networking research and implementation, the Joint Engineering Team (JET) recommends further identification and prioritization of these items followed by a proactive pursuit of these actions through its monthly meetings and other venues, as appropriate. This will particularly be pursued during its next two to four meetings.

3. Service Provisioning

The discussion on service provisioning focused on the need to securely allocate bandwidth and/or light paths within, or across, autonomous domains while efficiently sharing scarce network resources.

In the past, intra- and inter-domain provisioning was not implemented across network boundaries for various reasons, despite the existence of capable technologies like ATM. Today, modifications and extensions to existing boundary protocols (like BGP) would permit delineation of network boundaries and possibly permit secure inter-domain provisioning. Current inter-domain signaling protocols like MPLS may also provide a basis for intra- and inter-domain provisioning.

For inter-domain provisioning to be successful (manually or dynamically), all networks located along the end-to-end path must share a clearly defined common service offering. This issue is relevant to all levels of inter-domain hierarchical connections: RON-backbone, RON-RON, and RON-campus. The discussion eventually led to an agreement that the Global Lambda Integrated Facility (GLIF) is most likely the proper forum for the creation and maintenance of the service definitions. It was noted that the Quilt is also aware of this issue and interested in seeing resolution on it. Although no one at the workshop spoke authoritatively for the Quilt, those present believed it likely that the Quilt would follow GLIF's lead.

Two projects working in the intra- and inter-domain dynamic provisioning and routing area are the ESnet OSCARS project, and the Internet2 BRUW project. Both projects, being assisted by GEANT2 and CANARIE, have been coordinating their efforts and making progress in this area. Other JETnets interested in learning more about these projects or participating in the development process should contact Chin Guok at ESnet (chin@es.net), or Bob Riddle at Internet2/Abilene (bdr@internet2.edu). As appropriate, software will be made available to other JETnets.

Components of this topic that were identified by the Panels include:

- Intra- and inter-domain provisioning
- Resource sharing
- Dynamic and static networking
- Lambda allocation
- Prioritization and preemption of traffic
- MPLS tunnels
- User-controlled light paths
- Security, AAA
- Testing and inter-domain fault detection

4. Operations and Management

Operations and Management (O&M) provides the "care and feeding" of networks to enable the delivery of "production quality" services to end-users. O&M components that enable optical networking include monitoring network activity, understanding and managing traffic patterns, quick resolution of error conditions, and testing. It was noted that, as service offerings become more varied and complex (and are on multiple layers), fault isolation will become more complex. It will no longer be sufficient to "look for the break in the pipe," but will also be necessary to look for cases of mismatched descriptions—the leaks and obstructions.

Rick Summerhill (Internet2) stressed the importance of monitoring network performance and quickly and efficiently locating network faults. These operations could be provided by tools that allow the visualization of provisioned paths or virtual services on a network from a high-level, with the ability to "drill-down" to get specific performance statistics. These statistics are usually available from the command-line interface of the network devices, but a graphical tool might provide a better user experience. It was suggested that vendors should strive to provide better public APIs so that network operators can develop their own tools for monitoring. After the workshop concluded, it was learned that the NSF EDDY (End-to-end Diagnostic DiscoverY) project provides one such tool. EDDY collects data on application, network, system, environmental, and security events from multiple subsystems, normalizes them into a common event record, and disseminates them into a highly distributed diagnostic backplane where diagnosticians can perform both general and highly focused forensic analysis on a wide variety of faults and anomalies.

As it stands today, each JETnet has different levels of "acceptance testing" and methods it uses to monitor network health and performance. However, interagency cooperative network performance projects are taking place, and workshop participants suggested that cooperative sharing of network test equipment, particularly 10 Gigabit/s (Gbps) gear, could reduce the overall costs of network testing. Specific projects to test application-level performance over networks include Abilene's (among others) use of NetFlow data for flow performance analysis. In addition, Abilene/Internet2 is performing H.323 beacon testing, which injects H.323 streams into the network, then analyzes the results. New technologies like SFlow (Extreme and Force10 are two JETnet vendors implementing SFlow) have the capability of monitoring Layer 2 flows through a network device and can sample IPv6 as well as IPv4 traffic. Initiatives like the Quilt serve as forums for the exchange of information related to network acceptance testing.

Identified components of this topic include providing:
- Better views of what is happening across the network
- Better resource allocation and traffic management
- Capacity management
- Tools for fault isolation (looking for breaks, leaks, obstructions in the pipe)
- Measurement of end-to-end performance
- Tools for application performance monitoring
- Measurement for/as research
- Testing (10 Gigabit/s and beyond, lambda testing, export and flow analysis)

5. Cost Models

Participants, including NLR, Abilene/Internet2, NREN, ESnet, GEANT2, NISN and DREN discussed funding issues and the sharing of information regarding business models. To be effective, research funding needs to be sufficient and sustainable. However, in the current Federal climate, funds for research are particularly limited, and it is anticipated that in the future, users will be required to carry more of the financial burden.

The Quilt project was discussed as a model for demand aggregation, contract sharing, pricing information sharing and service sharing (including IP bearer service, shared NOC services and the swapping of light paths). IP transit sharing could help networks save on costs associated with providing commodity IP service and reduce the overhead and time spent on piecemeal peering.

Cost recovery is a concern of several networks, including GEANT2, which mentioned that while they experience marginal costs for light paths over an existing network, certain costs (e.g., transponders) are passed directly on to the participants and thus do not impact the financial situation of the network provider. This particular model encourages the use of light paths. For those providers that permit nonmembers to have their own light paths, in many cases, the fees are higher than those for members. This helps cover shared costs of the network not recovered by the charges for transponders for example.

The JETnets should expand their cooperation and collaboration to reduce costs and increase the effectiveness of the limited research resources. JETnets also need to work together to provide outreach to ensure that elected, appointed, and staff members of the legislative and executive branches are aware of the developments, improvements, and increasing importance of networks and network research—with particular emphasis on how these networks provide engines for economic development and contribute to maintaining U.S. competitive advantages.

With the growth of computational science, its reliance on high-performance networks, and its use by industry, continued improvements to the Internet are essential for economic development. Outreach activities might include conducting demonstrations and providing educational brochures. The recent President's Information Technology Advisory Committee report, "Computational Science: Ensuring America's Competitiveness," may help in this area. To ensure that the latest innovations, tools, and techniques are presented, the JETnets should coordinate with industry representatives where appropriate.

Identified components of this topic include:

- Funding
- Sharing of information
- Costs, cost recovery, pricing
- Stability of research funding in a time of limited Federal resources
- Security
- Educational literature

6. "MorphNet" Architecture

Most Federal JETnets carry some production traffic, as well as conduct research into new technologies and protocols. Workshop participants noted issues of potential conflicts in supporting both research and production activities over one network.

NLR, having been built on the "MorphNet" model (see: http://www.anl.gov/ECT/Public/research/morphnet.html) is described as "not a single network, but a set of facilities, capabilities, and services to build both experimental and production networks at various layers, allowing members to acquire dedicated (project-specific) facilities or shared (community-specific) facilities as appropriate." GEANT2's description is very similar. Some RONs also are built on the MorphNet model.

Workshop participants noted the need to migrate experimental research technologies onto production networks, but expressed concern over the lack of a well-defined process to do so. For example, the Department of Energy has two networks: UltraScienceNet and ESnet. How should technologies be migrated from USN to ESnet without disrupting existing production traffic? One approach that networks like GEANT2 have taken is to provide Layer 1 optical wavelengths to researchers, while operating a production Layer 2/3 network on other wavelengths. Another possible solution is for Layer 2 Ethernet VLANs to be provided to researchers while operational network business is operated on other VLANs at Layer 2 or 3.

Summary Report on "Service Definitions"

Jerry Sobieski, Chair

1. Overview

With the recent demand for, and provisioning of new Layer 1 and Layer 2 communications services, such as optical light paths and Ethernet framed service, comes the need for specific definitions of these services, so that providers can specify the exact services they are providing in response to customer requests.

Currently, there are many parameters left undefined when a user requests one of these new-styled services. For example, "100 Megabit/s Ethernet service, from point A to point B" could be interpreted in many ways:

- Half duplex, or full duplex?
- Tagged VLANs, or untagged?
- 1,500 byte frame size, or 4,470, or 9,216 bytes?
- Full bandwidth, or shared?

Common service definitions and formalized descriptions are needed for these service features, so that appropriate service options can be presented and understood by all parities—from the service provider all the way to the end consumer, as appropriate.

Architects and engineers from across the research and education networking community should sit down and discuss the formalization of these definitions. It is important to develop tangible specifications that customers can measure, and engineers from all appropriate parties can test to the service levels agreed upon during negotiation and contract phases.

A white paper discussing the justification and associated issues surrounding the topic, "Common Service Definitions" (Sobieski/Lehman) can be found at http://dragon.maxgigapop.net/public/Common-Services-Definition-v0.6.pdf or by contacting Jerry Sobieski (Mid-Atlantic Crossroads/University of Maryland) directly at jerrys@maxgigapop.net.

At the Global Lambda Integrated Facility (GLIF) meeting immediately following the ONT2 workshop, the "GLIF Control Plane" WG, chaired by Gigi Karmous-Edwards (MCNC), continued its work toward automating end-to-end connections over GLIF resources, with the ultimate goal of using control plane signaling and restoration. The WG heard first from Jerry Sobieski on "Common Service Definitions," then from Bill St. Arnaud (CANARIE/CA*net4) on the usefulness of Service Oriented Architectures (SOAs) and their potential role in automation, and finally from Matt Zekauskas (Internet2) on testing and verification of services.

It was realized that common service definitions may need to have more application type parameters as well (e.g., jitter and latency), and that many services will need mechanisms for the user to verify service delivery on an end-to-end basis. A task force led by Sobieski will continue to work towards developing the initial set of common service definitions by the end of 2005.

2. Actionable Objectives

2.1 Testbeds like **DRAGON should identify key services** they need to offer their users, **then develop engineering specifications** (including cost/price as appropriate), followed by encodings in appropriate formal languages, e.g., BNF or XML.

2.2 **Set up discussions on this topic at appropriate forums** within the community to inform each other and compare work in this area taking place within ITU, OIF, Quilt, Global Grid Forum (GGF), and GLIF, using open venues like iGRID, SC05, GGF and GLIF meetings.

2.3 The community should get started right away to **define and adopt an initial set of basic light path service definitions** by working within the GLIF framework.

2.4 The community should have a number of **Open Exchange Points up and running with services based on agreed common service definitions by September 2006**, and review progress in this area at ONT3.